Shifting Segments to Optimality

Stefan Felsner

Dedicated to Rolf H. Möhring

A rectangular dissection is a partition of a frame rectangle into rectangles, Figure 1 shows an example. Rectangular dissections are studied in various fields:

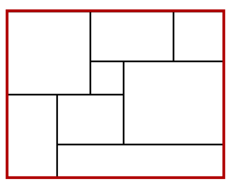

Figure 1: A rectangular dissection.

- architects look at them in the context of floorplan generation [11, 14],
- floorplaning is relevant for module placement in VLSI design [4, 21],
- in graph drawing rectangular dissections play a role in various representation models for planar graphs [12, 8],
- in cartography rectilinear dissections are studied as a special class of cartograms [13, 19, 17].

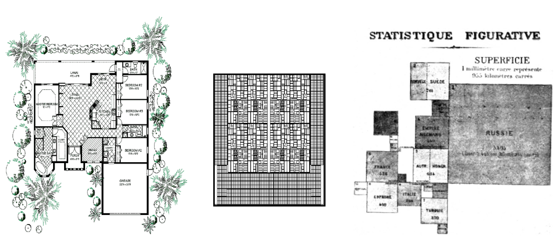

Figure 2:
Near rectangular dissections in applications.

Note that in the applications the areas of the rectangles of a dissection are relevant, actually in most cases these areas are prescribed. A rectangular dissection is *area-universal* if for any assignment of positive weights to the rectangles there is an equivalent dissection such that the areas of rectangles are equal to the given weights.

The key player in this article is Theorem 2 which characterizes area-universal rectangular dissections. In Section 2 we state the theorem and discuss some proofs and generalizations. Before getting there we present two problems which do not appear to have much in common. In Section 3 we show that both problems can be solved by clever applications of the theorem.

1 Two Problems

From now on a *dissection* shall be a rectangular dissection unless it is explicitly said. A dissection is *generic* if it has no cross, i.e., no point where four rectangles of the partition meet. A *segment* of a dissection is a maximal nondegenerate interval that belongs to the union of the boundaries of the rectangles. In general we disregard the four segments from the boundary frame, i.e., we only consider inner segments. Segments are either horizontal or vertical. The segments of a generic dissection are internally disjoint. Two dissections R and R' are *weakly equivalent* if there exist a bijection ϕ between their segments that preserves the orientation (horizontal/vertical) and such that segment s has an endpoint on segment t in R iff $\phi(s)$ has an endpoint on $\phi(t)$ in R'. A set P of points in $\mathbb{R}^2$ is *generic* if no two points from P have the same x or y coordinate.

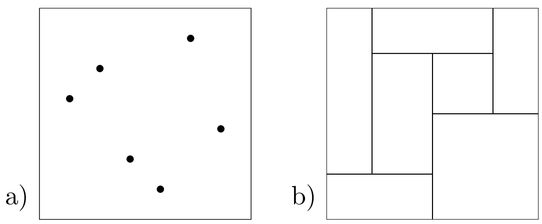

Figure 3: A generic set of six points and a generic dissection with six segments.

Let P be a set of n points in a rectangle frame F and let R be a generic dissection with n segments. A *cover map* from R to P is a dissection R' that is weakly equivalent to R and has outer rectangle F such that every segment of R' contains exactly one point from P. Figure 4 shows an example.

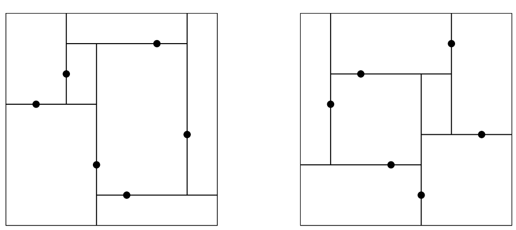

Figure 4: Two cover maps from the dissection of Fig. 3.b to the point set of Fig. 3.a.

Ackerman et al. [1] posed the following problem:

Problem 1 *Does a cover map from R to P exist for all pairs (R, P) where R is a generic dissection with n segments and P is a generic set of n points?*

The second problem is about *rectilinear duals* of planar graphs. In this drawing model the vertices are represented by simple rectilinear polygons, while edges are represented by side-contacts between the corresponding polygons, see Figure 5.

Now assume that positive weights $w(v)$ have been assigned to the vertices of the graph. A *rectilinear cartogram* is a rectilinear dual with the additional property that for all vertices the area of the polygon representing v equals $w(v)$. A relevant parameter measuring the complexity of a cartogram is the maximum number of sides of any polygon.

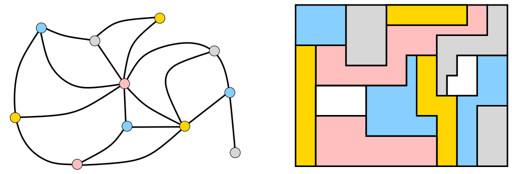

Figure 5: A graph with a rectilinear dual containing two white holes.

Problem 2 *What is the minimum number k such that any given planar triangulation with positive weights $w(v)$ admits a rectilinear cartogram with $\leq k$-sided polygons in a rectangular frame F of area $\sum_v w(v)$?*

A counting argument shows that the minimum k has to be at least 8. To see this note that a vertex of degree 3 enforces that at least one of the polygons representing the neighbors has a concave corner. Now let T^+ be a triangulation obtained from a triangulation T by stacking a new vertex of degree 3 into each face of T. If T has n vertices, then there are $2n-5$ new vertices in T^+. Together the n polygons representing the old vertices have at least $2n-5$ concave corners. If $n > 5$ there is at least one polygon with two concave corners, hence, with at least 8 sides.

2 The tool: area-universality of the weak equivalence class

The following theorem is a formalization of the statement given in the heading: *weak equivalence classes are area-universal.*

Theorem 1 *Let R be a dissection with rectangles $r_1, \ldots, r_{n+1}$, let the frame F be a rectangle and let $w : \{1, \ldots, n+1\} \to \mathbb{R}_+$ be a weight function with $\sum_i w(i) = \mathsf{area}(F)$. There exists a unique dissection R' contained in F that is weakly equivalent to R such that the area of the rectangle $\phi(r_i)$ in R' is $w(i)$.*

This theorem was first proven by Wimer et al. [21]. They take the width x_i and height y_i of rectangle r_i as variables and show that the system consisting of linear equations which correspond to left-to-right and bottom-to-top sequences of rectangles together with the non-linear equations $x_i y_i = w(i)$ has a unique solution. The theorem was rediscovered by Eppstein et al. [5]. They prove it with an argument based on "*invariance of domain*". Both proofs are purely existential though in [5, 6] it is noted that the solution can be computed by iteratively reducing the distance between the present weights and the intended weight vector $(w(i))_i$. Below we discuss an iterative approach and a corresponding proof based on 'air-pressure' in some detail. Before getting there, however, we introduce a class of dissections so that we can state the most important special case of the theorem.

Two dissections are *dual equivalent* if they have the same dual graph. In most applications we are interested in finding an appropriate member of the dual equivalence class. If two dissections are weakly equivalent they need not be dual equivalent, for example in Figure 4 the rectangle in the lower left corner has 4 neighbors in the left dissection but only 3 in the right dissection.

A segment s of a dissection is *one-sided* if s is the side of at least one of the rectangles, in other words all the segments that have an endpoint on s are on the same side of s. A dissection is *one-sided* if every segment of the dissection is one sided. The following observation was made in [6].

Proposition 1 *All dissections in the weak equivalence class of a one-sided dissection are dual equivalent.*

Together with Theorem 1 this yields the key theorem.

Theorem 2
One-sided dissections are area-universal.

With the following definitions we set the stage for a generalization of Theorem 1. Let $\mu : [0,1]^2 \to \mathrm{IR}_+$ be a density function on the unit square with mass 1, i.e., $\int_0^1\int_0^1 \mu(x,y)dxdy = 1$. We assume that μ is well behaved so that all the integrals we need exist and are positive. The *mass* of an axis aligned rectangle $r \subseteq [0,1]^2$ is defined as $m(r) = \iint_r \mu(x,y)dxdy$.

Theorem 3 *Let $\mu : [0,1]^2 \to \mathrm{IR}_+$ be a density function on the unit square. If R is a dissection with rectangles $r_1,\ldots,r_{n+1}$ and $w : \{1,\ldots,n+1\} \to \mathrm{IR}_+$ a positive weight function with $\sum_1^{n+1} w(i) = 1$ then there exists a unique dissection R' in the unit square that is weakly equivalent to R such that the mass of the rectangle $\phi(r_i)$ in R' is exactly $w(i)$.*

In [9] this theorem was proven with the air-pressure technique proposed by Izumi, Takahashi and Kajitani [10]. A short non-constructive proof of Theorem 3 was given by Schrenzenmaier [16, page 21], he adapted the proof of Theorem 1 from [5].

We describe the idea from [9]: Consider a realization of R in the unit square and compare the mass $m(r_i)$ to the intended mass $w(i)$. The quotient of these two values can be interpreted as the pressure inside the rectangle. Integrating this pressure along a side of the rectangle yields the force by which r_i is pushing against the segment that contains the side. The difference of pushing forces from both sides of a segment yields the effective force acting on the segment. The intuition is that shifting a segment in direction of the effective force yields a better balance of pressure in the rectangles. We show that iterating such improvement steps drives the realization of R towards a situation with $m(r_i) = w(i)$ for all i, i.e., the procedure converges towards the dissection R' whose existence we want to show.

Let $r_i = [x_l, x_r] \times [y_b, y_t]$ be a rectangle of R. The pressure $p(i)$ in r_i is $p(i) = \frac{w(i)}{m(r_i)}$. Let s be a segment of R and let r_i be one of the rectangles with a side in s. Let s be vertical with x-coordinate x_s and let $s \cap r_i$ span the interval $[y_b(i), y_t(i)]$. The (undirected) *force imposed on s by r_i* is the pressure $p(i)$ of r_i times the density dependent length of the intersection.

$$f(s,i) = p(i)\int_{y_b(i)}^{y_t(i)} \mu_{x_s}(y)dy.$$

The *force acting on s* is obtained as a sum of the directed forces imposed on s by incident rectangles.

$$f(s) = \sum_{r_i \text{ left of } s} f(s,i) - \sum_{r_i \text{ right of } s} f(s,i).$$

Symmetric definitions apply to horizontal segments.

Balance for rectangles and segments

Definition 1 *A segment s is in* balance *if $f(s) = 0$. A rectangle r_i is in* balance *if $m(r_i) = w(i)$, i.e., if $p(i) = 1$.*

Lemma 1 *All rectangles r_i of R are in balance if and only if all segments are in balance.*

Proof. We only show one direction. Since all rectangles are in balance we can eliminate the pressures from the definition of the $f(s, i)$. With this simplification we get for a vertical segment s

$$f(s) = \sum_{r_i \text{ left}} \int_{y_b(i)}^{y_t(i)} \mu_{x_s}(y)dy - \sum_{r_j \text{ right}} \int_{y_b(j)}^{y_t(j)} \mu_{x_s}(y)dy.$$

Hence $f(s) = M_s - M_s = 0$, where M_s is the integral of the fiber density μ_{x_s} along s. $\square$

Balancing segments and optimizing the entropy

Proposition 2 *If a segment s of R is unbalanced, then we can keep all the other segments at their position and shift s parallel to a position where it is in balance. The resulting dissection R' is weakly equivalent to R.*

Definition 2 *The* entropy *of a rectangle r_i of R is defined as $-w(i) \log p(i)$. The* entropy *of the dissection R is*

$$E = \sum_i -w(i) \log p(i)$$

The proof of Theorem 3 is in five steps:

(1) The entropy E is always nonpositive.

(2) $E = 0$ if and only if all rectangles r_i of R are in balance.

(3) Shifting an unbalanced segment s into its balance position increases the entropy.

(4) The process of repeatedly shifting unbalanced segments into their balance position makes R converge to a dissection R' such that the entropy of R' is zero.

(5) The solution is unique.

3 The solutions

Mapping segments on points

Let R be a generic dissection and S be a subset of the segments of R of size n and let P be a generic set of n points in a rectangle F. A *cover map* from (R, S) to P is a dissection R' with outer rectangle F that is weakly equivalent to R such that every segment in $S' = \phi(S)$ contains exactly one point from P. The following theorem answers our first problem.

Theorem 4 *If R is a generic dissection with a prescribed subset S of the segments of size n and P is a generic set of n points in a rectangle F, then there is a cover map R' from (R, S) to P.*

To be able to use Theorem 3 we first transform the point set P into a suitable density distribution $\mu = \mu_P$ inside F. This density is defined as the sum of a uniform distribution μ_1 with $\mu_1(q) = 1/\mathsf{area}(F)$ for all $q \in F$ and a distribution μ_2 that represents the points of P. Choose some $\Delta > 0$ such that for all $p, p' \in P$ we have $|x_p - x_{p'}| > 3\Delta$ and $|y_p - y_{p'}| > 3\Delta$, this is possible because P is generic. Define $\mu_2 = \sum_{p \in P} \mu_p$ where $\mu_p(q)$ takes the value $(\Delta^2\pi)^{-1}$ on the disk $D_\Delta(p)$ of radius Δ around p and value 0 for q outside of this disk.

For a density ν over F and a rectangle $r \subseteq F$ we let $\nu(r)$ be the integral of the density ν over r. Using this notation we can write $\mu_1(F) = 1$ and $\mu_p(F) = 1$ for all $p \in P$, hence the total mass of F is $\mu(F) = 1 + n$.

Next we transform the dissection R into a dissection R_S depending on the set S of segments that has to cover the points of P. To this end we replace every segment in S by a thin rectangle, see Figure 6. Let $\mathcal{S}$ be the set of new rectangles obtained from segments of S.

Define weights for the rectangles of R_S as follows. If R_S has r rectangles we define $w(r) = 1 + 1/r$ if $r \in \mathcal{S}$ and $w(r) = 1/r$ for all the rectangles of R_S that came from rectangles of R. Note that the total weight, $\sum_r w(r) = 1 + n$, is in correspondence to the total mass $\mu(R)$.

The data R with μ and R_S with w constitute, up to scaling of R and w, a set of inputs for Theo-

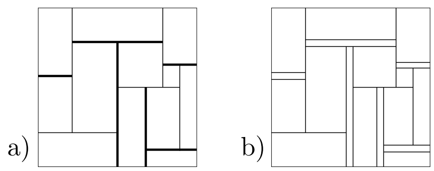

Figure 6: Dissections R with as prescribed subset S of segments (bold and gray) and the dissection R_S obtained by doubling the segments of S.

rem 3. From the conclusion of the theorem we obtain a dissection R'_S weakly equivalent to R_S such that $m(r) = \iint_r \mu(x,y)dxdy = w(r)$ for all rectangles r of R'_S.

The definition of the weight function w and the density μ is so that R'_S should be close to a cover map from (R,S) to P: Only the rectangles $r \in \mathcal{S}$ that have been constructed by inflating segments may contain a disk $D_\Delta(p)$ and each of these rectangles may contain at most one of the disks. This suggests a correspondence $\mathcal{S} \leftrightarrow P$. However, a rectangle $r \in \mathcal{S}$ can use parts of several discs to accumulate mass. To find a correspondence between $\mathcal{S}$ and P we define a bipartite graph G whose vertices are the points in P and the rectangles in $\mathcal{S}$:

- A pair (p,r) is an edge of G iff $r \cap D_\Delta(p) \neq \emptyset$ in R'_S.

The proof of the theorem is completed by proving two claims:

- G admits a perfect matching.
- From R'_S and a perfect matching M in G we can produce a dissection R' that realizes the cover map from (R,S) to P.

For the first of the claims we check Hall's matching condition: Consider a subset A of $\mathcal{S}$. Since R_S is realizing the prescribed weights we have $m(A) = \mu(A) = \sum_{r\in A}\mu(r) = \sum_{r\in A} w(r) = |A|(1+1/r)$. Since $\mu_1(A) < 1$ and $\mu_p(A) \leq 1$ for all $p \in P$ there must be at least $|A|$ points $p \in P$ with $\mu_p(A) > 0$, these are the points that have an edge to a rectangle from A in G. We have thus shown that every $A \subset \mathcal{S}$ is incident to at least $|A|$ points in G, hence, there is an injective mapping $\alpha : \mathcal{S} \to P$ such that $r \cap D_\Delta(\alpha(r)) \neq \emptyset$ in R'_S for all $r \in \mathcal{S}$.

Given the matching α the construction of the dissection R' that realizes the cover map from (R,S) to P is completed in three steps as indicated in Figure 7.

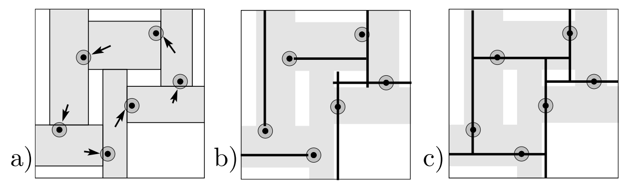

Figure 7: a) A solution R'_S with a matching indicated by the arrows. b) Segments $s \in S$ shifted to the corresponding points. c) Small final adjustments (clipping and enlarging) yield R'.

Cartograms with optimal complexity

In a series of papers the complexity of polygons has been reduced from the initial 40 to 34 then 12 and 10. Finally, in [3] the following optimal result was obtained.

Theorem 5 *Every planar triangulation admits an area-universal rectilinear cartogram with* $\leq$8*-sided polygons.*

The construction is fairly easy with the right tools at hand. First we need a Schnyder wood of the input triangulation G. Let a_1, a_2, a_3 be the outer vertices of G, an orientation and coloring of the inner edges with with 3 colors (we identify colors (1,2,3) with (red,green,blue)) is a *Schnyder wood* if:

(1) All edges incident to an outer vertex a_i are in-edges and colored i.

(2) Every inner vertex v has three outgoing edges colored red, green and blue in clockwise order. All the incoming edges in an interval between two outgoing edges are colored with the third color, see Figure 8.

These structures were defined by Schnyder in [15], there it is shown that every triangulation admits a Schnyder wood. Moreover, if T_i is the set of oriented edges of color i and T_i^{-1} is the same set with reversed orientations, then it holds that $T_1 \cup T_2^{-1} \cup T_3^{-1}$ is

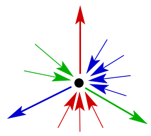

Figure 8: Schnyder's edge coloring rule.

acyclic. This property can be used to show that every triangulation has a contact representation with internally disjoint $\perp$ shapes. Figure 9 shows an example.

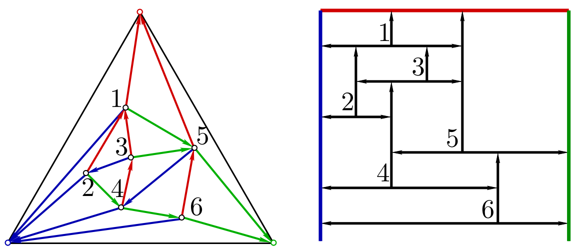

Figure 9: A triangulation with a Schnyder wood and a $\perp$-representation

The $\perp$-representation can be viewed as a rectangular dissection. Now replace every segment of this dissection by a thin rectangle. This yields a one-sided dissection R_G, see Figure 10(left). With a vertex v of G we associate the polygon P_v formed as the union of the two rectangles that were obtained from the two segments of the $\perp$ shape representing v together with the two rectangles that have parts of the horizontal segment of this $\perp$ as bottom side, see Figure 10(right).

It is easily checked that the polygons P_v have at most 8 corners, hence, at most 8 sides. Given a set of weights $w : V \to \mathbb{R}^+$ we can arbitrarily break $w(v)$ into four positive values and assign these to the rectangles constituting P_v. Since the dissection R_G is one-sided and, hence, area-universal there is a realization of the dissection where the area of $P(v)$ equals $w(v)$.

From the thesis of Torsten Ueckerdt [20] we borrow our last figure which shows a cartogram displaying real data. The cartogram was computed with the method of this section.

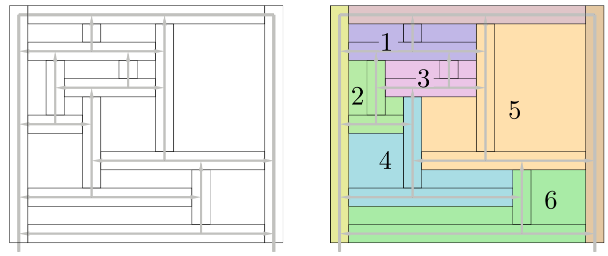

Figure 10: The one-sided dissection resulting from the $\perp$-representation of Figure 9.

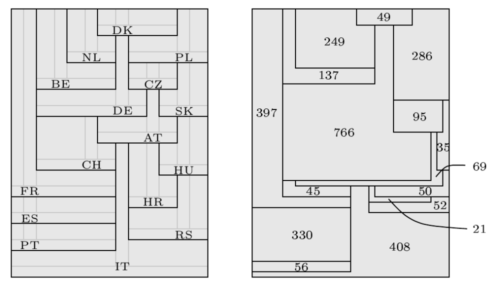

Figure 11: Central European states represented by polygons with at most 8 sides (left). A cartogram where the areas are proportional to the emission of CO_2 in 2009 (right).

4 Background and additional problems

The two main problems regarding area-universal rectangular dissections are the following:

- Given R and w, is it possible to effectively compute the weakly equivalent dissection R' realizing the weights? (effective Theorem 1)
- Is it possible to effectively recognize graphs that admit a one-sided dual.

Beside this it would be very interesting to identify further instances of area-universality. Two such instances are straight line drawings of 3-regular planar graphs with prescribed face areas (Thomassen [18]) and straight line drawings of grids with prescribed face areas, a.k.a. table cartograms (Evan et al. [7]).

The conjecture of Ackerman et al. [1] regarding mappings of segments to point sets was motivated by the study of the function $Z(P)$ counting rectangulations of a generic point set P. Combining results from [9](lower) and [1](upper) we know that $Z(P)$ is in $\Theta(8^{n+1}/(n+1)^4)$ and in $O(20^n/n^4)$. The lower bound is tight for some sets P, to improve the upper bound remains an intriguing problem.

The construction of area-universal rectilinear cartograms with $\leq$8-sided polygons is from Alam et al. [3]. As already noted the construction based on our key theorem is not known to be effective. Effective constructions of cartograms with $\leq$8-sided polygons are known for certain Hamiltonian triangulations [3] and with $\leq$10-sided polygons for general triangulations [2] Can cartograms with $\leq$8-sided polygons be constructed effectively for general triangulations?

Recognition of planar graphs which admit rectangular cartograms or cartograms with $\leq$6-sided polygons is also wide open.

References

[1] E. Ackerman, G. Barequet, and R. Y. Pinter, *On the number of rectangulations of a planar point set*, J. Combin. Theory Ser. A, 113 (2006), 1072–1091.

[2] M. J. Alam, T. Biedl, S. Felsner, A. Gerasch, M. Kaufmann, and S. G. Kobourov, *Linear-time algorithms for rectilinear hole-free proportional contact representations*, Algorithmica, 67 (2013), 3–22.

[3] M. J. Alam, T. Biedl, S. Felsner, M. Kaufmann, S. G. Kobourov, and T. Ueckerdt, *Computing cartograms with optimal complexity*, Discr. and Comput. Geom., 50 (2013), 784–810.

[4] P. Dasgupta and S. Sur-Kolay, *Slicible rectangular graphs and their optimal floorplans*, ACM Trans. Design Autom. Electr. Syst., 6 (2001), 447–470.

[5] D. Eppstein, E. Mumford, B. Speckmann, and K. Verbeek, *Area-universal rectangular layouts*, 2009, 19 pages. arXiv:0901.3924.

[6] D. Eppstein, E. Mumford, B. Speckmann, and K. Verbeek, *Area-universal and constrained rectangular layouts*, SIAM J. Computing, 41 (2012), 537–564.

[7] W. Evans, S. Felsner, M. Kaufmann, S. G. Kobourov, D. Mondal, R. I. Nishat, and K. Verbeek, *Table cartograms*, in Proc. ESA, vol. 8242 of LNCS, Springer, 2013, 421–432.

[8] S. Felsner, *Rectangle and square representations of planar graphs*, in Thirty Essays in Geometric Graph Theory, J. Pach, ed., Springer, 2013, 213–248.

[9] S. Felsner, *Exploiting air-pressure to map floorplans on point sets*, J. Graph Alg. and Appl., (2014), 233–252.

[10] T. Izumi, A. Takahashi, and Y. Kajitani, *Air-pressure model and fast algorithms for zero-wasted-area layout of general floorplan*, IEICE Trans. Fundam. Electron., Commun. and Comp. Sci., E81–A (1998), 857–865.

[11] K. Nassar, *New advances in the automated architectural space plan layout problem*, in Proceedings Computing in Civil and Building Engineering, 2010, 9 pages. http://www.engineering.nottingham.ac.uk/icccbe/proceedings/pdf/pf193.pdf.

[12] T. Nishizeki and M. S. Rahman, *Planar Graph Drawing*, vol. 12 of Lecture Notes Series on Computing, World Scientific, 2004.

[13] E. Raisz, *The rectangular statistical cartogram*, Geographical Review, 24 (1934), 292–296.

[14] I. Rinsma, *Non-existence of a certain rectangular floorplan with specified area and adjacency*, Environment and Planning, 14 (1987), 163–166.

[15] W. Schnyder, *Planar graphs and poset dimension*, Order, 5 (1989), 323–343.

[16] H. Schrenzenmaier, *Ein Luftdruckparadigma zur Optimierung von Zerlegungen.* Bachelor's Thesis, TU Berlin, 2013.

[17] W. Team, *Worldmapper, the world as you've never seen it before.* http://www.worldmapper.org.

[18] C. Thomassen, *Plane cubic graphs with prescribed face areas*, Comb., Prob. & Comp., 1 (1992), 371–381.

[19] W. Tobler, *Thirty five years of computer cartograms*, Annal. Assoc. Am. Geogr., 94 (2004), 58–73.

[20] T. Ueckerdt, *Geometric representations of graphs with low polygonal complexity.* Dissertation, TU Berlin, 2011.

[21] S. Wimer, I. Koren, and I. Cederbaum, *Floorplans, planar graphs, and layouts*, IEEE Transactions on Circuits and Systems, 35 (1988), 267 –278.